

UNITED STATES PATENT APPLICATION

of

Chad LESTER

and

Eric V EACH

for

ENCODING AD AND/OR AD SERVING INFORMATION IN A CLICK URL

Expr ss Mail Lab | No. EL997188822US

ENCODING AD AND/OR AD SERVING INFORMATION IN A CLICK URL

§ 0. RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Application Serial No.

5 60/442,209, titled "METHOD FOR ENCODING ADVERTISEMENT INFORMATION IN A CLICK URL," filed on January 24, 2003 and listing Eric Hugh Veach and Chad Lester as the inventors. That application is expressly incorporated herein by reference. The present invention is not limited to any requirements described in the provisional application.

10

§ 1. BACKGROUND OF THE INVENTION

§ 1.1 FIELD OF THE INVENTION

The present invention concerns advertising. In particular, the present invention concerns tracking ad selection and ad performance.

15

§ 1.2 BACKGROUND INFORMATION

Advertising using traditional media, such as television, radio, newspapers and magazines, is well known. Unfortunately, even when armed with demographic studies and entirely reasonable assumptions about the typical audience of various media outlets, advertisers recognize that much of their ad budget is simply wasted. Moreover, it may be very difficult to identify and eliminate such waste.

20 Recently, advertising over more interactive media has become popular. For example, as the number of people using the Internet has exploded, advertisers have come to appreciate media and services offered over the Internet as a potentially 25 powerful way to advertise.

Advertisers have developed several strategies in an attempt to maximize the 30 value of such advertising. In one strategy, advertisers use popular presences or means for providing interactive media or services (referred to as "Websites" in the specification without loss of generality) as conduits to reach a large audience. Using this first approach, an advertiser may place ads on the home page of the New York Times Website, or the USA Today Website, for example. In another strategy, an advertiser

may attempt to target its ads to narrower niche audiences, thereby increasing the likelihood of a positive response by the audience. For example, an agency promoting tourism in the Costa Rican rainforest might place ads on the ecotourism-travel subdirectory of the Yahoo Website. An advertiser will normally determine such

5 targeting manually.

Regardless of the strategy, Website-based ads (also referred to as "Web ads") are often presented to their advertising audience in the form of "banner ads" – i.e., a rectangular box that includes graphic components. When a member of the advertising audience (referred to as a "viewer" or "user" in the specification without loss of

10 generality) selects one of these banner ads by clicking on it, embedded hypertext links typically direct the viewer to a page on the advertiser's Website (referred to as an "ad landing page" or simply a "landing page"). This process, wherein the viewer selects an ad, is commonly referred to as a "click-through" ("Click-through" is intended to cover any user selection.). The ratio of the number of click-throughs to the number of impressions

15 of the ad (i.e., the number of times an ad is rendered) is commonly referred to as the "click-through rate" of the ad.

A "conversion" is said to occur when a user consummates a transaction related to a previously served ad. What constitutes a conversion may vary from case to case and can be determined in a variety of ways. For example, it may be the case that a

20 conversion occurs when a user clicks on an ad, is referred to the advertiser's web page, and consummates a purchase there before leaving that Web page. Alternatively, a conversion may be defined as a user being shown an ad, and making a purchase on the advertiser's Web page within a predetermined time (e.g., seven days). In yet another alternative, a conversion may be defined by an advertiser to be any

25 measurable/observable user action such as, for example, downloading a white paper, navigating to at least a given depth of a Website, viewing at least a certain number of Web pages, spending at least a predetermined amount of time on a Website or Web page, registering with a Website, etc. Often, if user actions don't indicate a consummated purchase, they may indicate a sales lead, although user actions

30 constituting a conversion are not limited to this. Indeed, many other definitions of what constitutes a conversion are possible. The ratio of the number of conversions to the

number of impressions of the ad (i.e., the number of times an ad is displayed) is commonly referred to as the conversion rate. If a conversion is defined to be able to occur within a predetermined time since the serving of an ad, one possible definition of the conversion rate might only consider ads that have been served more than the
5 predetermined time in the past.

The serving of ads can be priced in various ways. Common examples include cost per impression (CPI) and cost per action (CPA), where “action” refers to some desired activity of the user such as selecting (or clicking) the ad, purchasing a product, signing up for a service, or some other “conversion.” Initially, online advertising was
10 generally priced per impression (e.g., per thousand impressions). However, pricing schemes which more closely align online advertising costs with the desired results of the advertiser (e.g., having a user visit the advertiser’s Website, having a user consummate a commercial transaction on the advertiser’s Website, etc.) have become
15 more popular. To track user selections (clicks), the links in some ads may include links that first direct the user’s content rendering application (e.g., a browser) to one or more intermediate resources or servers (e.g., using one or more URLs). These servers associated with intermediate URLs (simply referred to as “intermediate URL servers” without loss of generality) can direct the user’s content rendering application (e.g., a browser) to a server that can record the selection (click event) before forwarding the
20 user’s content rendering application to a subsequent server (which may be the server storing the ad landing page). Note that in addition to tracking selections for billing purposes, selections can be used to derive ad performance metrics. Such metrics may be used for various purposes such as positioning ads or applying relative rendering attributes, ad campaign assistance, governing the serving of ads, etc.

25 As can be appreciated from the foregoing, there is a need for methods and apparatus that allow clicks to be tracked. Moreover, such tracking should support performance monitoring, perhaps supporting rich levels of performance monitoring.

§ 2. SUMMARY OF THE INVENTION

30 The present invention helps to track ad selections, such as ad clicks. The present invention may also help to track rich levels of ad performance. The present

invention may do so by encoding one or more ad properties, such as an identity of the ad; an identity of the advertiser; a time the ad was served; a time the ad was rendered; a rendering attribute of the ad; a position of the ad within a Web page; a price that the advertiser will be charged for the impression; a price that the advertiser will be charged

5 for a click; a price that the advertiser will be charged for a conversion; an identity of the server that chose the ad; search conditions that generated the page with which the ad was rendered; a next intermediate URL; a final destination URL; an identity of the ad creative; a topic of the content or other information about the content with which the ad was served (e.g. one or more concepts associated with the Web page on which the 10 advertisement was displayed); an identity of the content with which the ad was served; information about other advertisements that were rendered along with the ad; a geolocation to which the ad was served; and user profile information of the user to which the ad was served, various other ad serving parameters, etc., in a click URL of the ad.

15 An intermediate URL server may be used to decode such ad properties when the ad is selected.

§ 3. BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a high-level diagram showing parties or entities that can interact with 20 an advertising system.

Figure 2 illustrates an environment in which advertisers can target their ads on search results pages generated by a search engine, documents served by content servers, and/or e-mail.

25 Figure 3 is a bubble diagram of operations that may be performed consistent with the present invention.

Figure 4 is a flow diagram of an exemplary method that may be used to provide encoded ad property information in a click URL in a manner consistent with the present invention.

30 Figure 5 is a flow diagram of an exemplary method that may be used to encode variable length information in a manner consistent with the present invention.

Figure 6 is a block diagram of apparatus that may be used to perform at least some of the various operations that may be used, and store at least some of the information that may be used and/or generated, in a manner consistent with the present invention.

5

§ 4. DETAILED DESCRIPTION

The present invention may involve novel methods, apparatus, message formats, and/or data structures for encoding ad property information and providing such encoded information in a click URL. The following description is presented to enable one skilled in the art to make and use the invention, and is provided in the context of particular applications and their requirements. Various modifications to the disclosed embodiments will be apparent to those skilled in the art, and the general principles set forth below may be applied to other embodiments and applications. Thus, the present invention is not intended to be limited to the embodiments shown and the inventors regard their invention as any patentable subject matter described.

In the following, environments in which, or with which, the present invention may operate are described in § 4.1. Then, exemplary embodiments of the present invention are described in § 4.2. Finally, some conclusions regarding the present invention are set forth in § 4.3.

20

§ 4.1 ENVIRONMENTS IN WHICH, OR WITH WHICH, THE PRESENT INVENTION MAY OPERATE

§ 4.1.1 EXEMPLARY ADVERTISING ENVIRONMENT

25 Figure 1 is a diagram of an advertising environment. The environment may include an ad entry, maintenance and delivery system (simply referred to an ad server) 120. An exemplary ad server 120 is described in U.S. Patent Application Serial No. 10/610,311 (incorporated herein by reference), titled "PROMOTING AND/OR DEMOTING AN ADVERTISEMENT FROM AN ADVERTISING SPOT OF ONE TYPE 30 TO AN ADVERTISING SPOT OF ANOTHER TYPE," filed on June 30, 2003 and listing Chad Lester as the inventor, with reference to Figure 2. Advertisers 110 may directly, or indirectly, enter, maintain, and track ad information in the system 120. The ads may

be in the form of graphical ads such as so-called banner ads, text only ads, image ads, audio ads, video ads, ads combining one or more of any of such components, etc. The ads may also include embedded information, such as a link, and/or machine executable instructions. Ad consumers 130 may submit requests for ads to, accept ads responsive
5 to their request from, and provide usage information to, the system 120. An entity other than an ad consumer 130 may initiate a request for ads. Although not shown, other entities may provide usage information (e.g., whether or not a conversion or click-through related to the ad occurred) to the system 120. This usage information may include measured or observed user behavior related to ads that have been served. Ad
10 property information may also be provided to the system. As will be discussed below, such information may be encoded by the ad system 120 (e.g., when the ad is served) and gathered by an intermediate URL server.

Figure 2 illustrates an environment 200 in which the present invention may be used. A user device (also referred to as a "client" or "client device") 250 may include a
15 content rendering facility such as a browser facility (such as the Opera browser, the Explorer browser from Microsoft, or the Navigator browser from AOL/Time Warner), an e-mail facility (e.g., Outlook from Microsoft), etc. A search engine 220 may permit user devices 250 to search collections of documents (e.g., Web pages). A content server 210 may permit user devices 250 to access documents. An e-mail server (e.g., Hotmail
20 from Microsoft Network, Yahoo Mail, etc.) 240 may be used to provide e-mail functionality to user devices 250. An ad server 210 may be used to serve ads to user devices 250. The ads may be served in association with search results provided by the search engine 220, content provided by the content server 230, and/or e-mail supported by the e-mail server 240 and/or user device e-mail facilities.

25 Thus, one example of an ad consumer 130 is a general content server 230 that receives requests for documents (e.g., articles, discussion threads, music, video, graphics, search results, Web page listings, etc.), and retrieves the requested document in response to, or otherwise services, the request. The content server may submit a request for ads to the ad server 120/210. Such an ad request may include a number of
30 ads desired. The ad request may also include document request information. This information may include the document itself (e.g., page), a category or topic

corresponding to the content of the document or the document request (e.g., arts, business, computers, arts-movies, arts-music, etc.), part or all of the document request, content age, content type (e.g., text, graphics, video, audio, mixed media, etc.), geolocation information, document information, etc.

- 5 The content server 230 may combine the requested document with one or more of the advertisements provided by the ad server 120/210. This combined information including the document content and advertisement(s) is then forwarded towards the end user device 250 that requested the document, for presentation to the user. Exemplary content-relevant ad servers are described in U.S. Patent Application Serial No. 10/375,900 (incorporated herein by reference), titled "SERVING ADVERTISEMENTS BASED ON CONTENT," filed on February 26, 2003 and listing Darrell Anderson, Paul Bucheit, Alexander Paul Carobus, and Yingwei Cui as inventors. Finally, the content server 230 may transmit information about the ads and how, when, and/or where the ads are to be rendered, and/or ad performance (e.g., position, click-through or not, impression time, impression date, size, conversion or not, etc.) back to the ad server 120/210. Alternatively, or in addition, such information may be provided back to the ad server 120/210 by some other means. For example, such information may be encoded and provided in a click URL provided with the ad and decoded by an intermediate URL server.
- 10 Another example of an ad consumer 130 is the search engine 220. A search engine 220 may receive queries for search results. In response, the search engine may retrieve relevant search results (e.g., from an index of Web pages). An exemplary search engine is described in the article S. Brin and L. Page, "The Anatomy of a Large-Scale Hypertextual Search Engine," Seventh International World Wide Web Conference, Brisbane, Australia and in U.S. Patent No. 6,285,999 (both incorporated herein by reference). Such search results may include, for example, lists of Web page titles, snippets of text extracted from those Web pages, and hypertext links to those Web pages, and may be grouped into a predetermined number of (e.g., ten) search results.
- 15 The search engine 220 may submit a request for ads to the ad server 120/210. The request may include a number of ads desired. This number may depend on the
- 20 Another example of an ad consumer 130 is the search engine 220. A search engine 220 may receive queries for search results. In response, the search engine may retrieve relevant search results (e.g., from an index of Web pages). An exemplary search engine is described in the article S. Brin and L. Page, "The Anatomy of a Large-Scale Hypertextual Search Engine," Seventh International World Wide Web Conference, Brisbane, Australia and in U.S. Patent No. 6,285,999 (both incorporated herein by reference). Such search results may include, for example, lists of Web page titles, snippets of text extracted from those Web pages, and hypertext links to those Web pages, and may be grouped into a predetermined number of (e.g., ten) search results.
- 25 The search engine 220 may submit a request for ads to the ad server 120/210. The request may include a number of ads desired. This number may depend on the
- 30 Another example of an ad consumer 130 is the search engine 220. A search engine 220 may receive queries for search results. In response, the search engine may retrieve relevant search results (e.g., from an index of Web pages). An exemplary search engine is described in the article S. Brin and L. Page, "The Anatomy of a Large-Scale Hypertextual Search Engine," Seventh International World Wide Web Conference, Brisbane, Australia and in U.S. Patent No. 6,285,999 (both incorporated herein by reference). Such search results may include, for example, lists of Web page titles, snippets of text extracted from those Web pages, and hypertext links to those Web pages, and may be grouped into a predetermined number of (e.g., ten) search results.

search results, the amount of screen or page space occupied by the search results, the size and shape of the ads, etc. In one embodiment, the number of desired ads will be from one to ten, and preferably from three to five. The request for ads may also include the query (as entered or parsed), information based on the query (such as geolocation

5 information, whether the query came from an affiliate and an identifier of such an affiliate), and/or information associated with, or based on, the search results. Such information may include, for example, identifiers related to the search results (e.g., document identifiers or "docIDs"), scores related to the search results (e.g., information retrieval ("IR") scores such as dot products of feature vectors corresponding to a query 10 and a document, Page Rank scores, and/or combinations of IR scores and Page Rank scores), snippets of text extracted from identified documents (e.g., Web pages), full text of identified documents, topics of identified documents, feature vectors of identified documents, etc.

15 The search engine 220 may combine the search results with one or more of the advertisements provided by the ad server 120/210. This combined information including the search results and advertisement(s) is then forwarded towards the user that submitted the search, for presentation to the user. Preferably, the search results are maintained as distinct from the ads, so as not to confuse the user between paid advertisements and presumably neutral search results.

20 Finally, the search engine 220 may transmit information about the ad and when, where, and/or how the ad was to be rendered, and/or ad performance (e.g., position, click-through or not, impression time, impression date, size, conversion or not, etc.) back to the ad server 120/210. Alternatively, or in addition, such information may be provided back to the ad server 120/210 by some other means. For example, such 25 information may be encoded and provided in a click URL provided with the ad and decoded by an intermediate URL server.

Finally, as discussed in U.S. Patent Application Serial No. 10/452,830 (incorporated herein by reference), titled "SERVING ADVERTISEMENTS USING INFORMATION ASSOCIATED WITH E-MAIL," filed on June 2, 2003 and listing Jeffrey 30 A. Dean, Georges R. Harik and Paul Bucheit as inventors, the e-mail server 240 may be thought of, generally, as a content server in which a document served is simply an

e-mail. Further, e-mail applications (such as Microsoft Outlook for example) may be used to send and/or receive e-mail. Therefore, an e-mail server 240 or application may be thought of as an ad consumer 130. Thus, e-mails may be thought of as documents, and targeted ads may be served in association with such documents. For example, one 5 or more ads may be served in, under, over, or otherwise in association with an e-mail.

§ 4.1.3 DEFINITIONS

Online ads, such as those used in the exemplary systems described above with reference to Figures 1 and 2, or any other system, may have various intrinsic features.

10 Such features may be specified by an application and/or an advertiser. These features are referred to as "ad features" below. For example, in the case of a text ad, ad features may include a title line, ad text, and an embedded link. In the case of an image ad, ad features may include images, executable code, and an embedded link. Depending on the type of online ad, ad features may include one or more of the 15 following: text, a link, an audio file, a video file, an image file, executable code, embedded information, etc.

When an online ad is served, one or more parameters may be used to describe how, when, and/or where the ad was served. These parameters are referred to as "serving parameters" below. Serving parameters may include, for example, one or 20 more of the following: features of (including information on) a page on which the ad was served, a search query or search results associated with the serving of the ad, a user characteristic (e.g., their geolocation, the language used by the user, the type of browser used, previous page views, previous behavior), a host or affiliate site (e.g., America Online, Google, Yahoo) that initiated the request, an absolute position of the ad 25 on the page on which it was served, a position (spatial or temporal) of the ad relative to other ads served, an absolute size of the ad, a size of the ad relative to other ads, a color of the ad, a number of other ads served, types of other ads served, time of day served, time of week served, time of year served, etc. Naturally, there are other serving parameters that may be used in the context of the invention.

30 Although serving parameters may be extrinsic to ad features, they may be associated with an ad as serving conditions or constraints. When used as serving

conditions or constraints, such serving parameters are referred to simply as “serving constraints” (or “targeting criteria”). For example, in some systems, an advertiser may be able to target the serving of its ad by specifying that it is only to be served on weekdays, no lower than a certain position, only to users in a certain geolocation, etc.

5 As another example, in some systems, an advertiser may specify that its ad is to be served only if a page or search query includes certain keywords or phrases. As yet another example, in some systems, an advertiser may specify that its ad is to be served only if a document being served includes certain topics or concepts, or falls under a particular cluster or clusters, or some other classification or classifications.

10 “Geolocation information” may include information specifying one or more of one or more countries, one or more (inter-country) regions, one or more states, one or more metro areas, one or more cities, one or more towns, one or more boroughs, one or more areas with common zip codes, one or more areas with common telephone area codes, one or more areas served by common cable head end stations, one or more 15 areas served by common network access points or nodes, etc. It may include latitude and/or longitude, or a range thereof. It may include information, such as an IP address, from which a user location can be estimated.

20 “Ad information” may include any combination of ad features, ad serving constraints, information derivable from ad features or ad serving constraints (referred to as “ad derived information”), and/or information related to the ad (referred to as “ad related information”), as well as an extension of such information (e.g., information derived from ad related information).

“Ad properties” may include ad information, such as one or more ad features, and/or one or more serving parameters.

25 A “document” is to be broadly interpreted to include any machine-readable and machine-storable work product. A document may be a file, a combination of files, one or more files with embedded links to other files, etc. The files may be of any type, such as text, audio, image, video, etc. Parts of a document to be rendered to an end user can be thought of as “content” of the document. A document may include “structured 30 data” containing both content (words, pictures, etc.) and some indication of the meaning of that content (for example, e-mail fields and associated data, HTML tags and

associated data, etc.) Ad spots in the document may be defined by embedded information or instructions. In the context of the Internet, a common document is a Web page. Web pages often include content and may include embedded information (such as meta information, hyperlinks, etc.) and/or embedded instructions (such as Javascript, etc.). In many cases, a document has a unique, addressable, storage location and can therefore be uniquely identified by this addressable location. A universal resource locator (URL) is a unique address used to access information on the Internet.

“Document information” may include any information included in the document, information derivable from information included in the document (referred to as “document derived information”), and/or information related to the document (referred to as “document related information”), as well as an extensions of such information (e.g., information derived from related information). An example of document derived information is a classification based on textual content of a document. Examples of document related information include document information from other documents with links to the instant document, as well as document information from other documents to which the instant document links.

Content from a document may be rendered on a “content rendering application or device”. Examples of content rendering applications include an Internet browser (e.g., Explorer or Netscape), a media player (e.g., an MP3 player, a Realnetworks streaming audio file player, etc.), a viewer (e.g., an Adobe Acrobat pdf reader), etc.

A “content owner” is a person or entity that has some property right in the content of a document. A content owner may be an author of the content. In addition, or alternatively, a content owner may have rights to reproduce the content, rights to prepare derivative works of the content, rights to display or perform the content publicly, and/or other proscribed rights in the content. Although a content server might be a content owner in the content of the documents it serves, this is not necessary.

“User information” may include user behavior information and/or user profile information.

“E-mail information” may include any information included in an e-mail (also referred to as “internal e-mail information”), information derivable from information included in the e-mail and/or information related to the e-mail, as well as extensions of

such information (e.g., information derived from related information). An example of information derived from e-mail information is information extracted or otherwise derived from search results returned in response to a search query composed of terms extracted from an e-mail subject line. Examples of information related to e-mail 5 information include e-mail information about one or more other e-mails sent by the same sender of a given e-mail, or user information about an e-mail recipient. Information derived from or related to e-mail information may be referred to as "external e-mail information."

Various exemplary embodiments of the present invention are now described in 10 § 4.2.

§ 4.2 EXEMPLARY EMBODIMENTS

The present invention may help to track ad selections and various types of ad performance. The present invention may do so by encoding one or more ad properties. 15 Figure 3 is a bubble diagram of operations that may be performed consistent with the present invention. In the exemplary environment 300 shown, an ad server 310 may be used to provide ads or ad information 314 to content rendering (e.g., browser) operations 322 on a user device 320. Although the following description discusses browsers and Web pages, the present invention is applicable to other content rendering 20 applications or devices and other documents. When a user selects an ad, the browser 322 renders (goes to) a landing page 342 specified in the ad information 314. Thus, advertisements 314 displayed (or otherwise rendered) on the Web may include a link, which when followed, will ultimately lead the browser 322 to a Web page 342 of (the 25 advertiser's) web site server 340. However, this link may first direct the browser 322 to one or more intermediate URLs. These intermediate URLs can direct the browser 322 to an intermediate URL server 330. At the intermediate URL server 330, action event recording operations 332 can record the user selection (e.g., click event) and browser 30 redirect operations 334 can forward the browser 322 to a subsequent URL, such as that of the landing page 342 for example. The various servers 310, 320, 330, and 340 can communicate with one another via one or more networks 350, such as the Internet for example. As indicated by dashed line 360, the ad server 310 and the intermediate URL

server 330 may reside at the same location, and/or may be under control of the same entity.

Ad property encoding operations 312 may be used to generate an intermediate URL that encodes one or more properties of the ad, such as an identity of the ad; an identity of the advertiser; a time (local at server or local at user device) the ad was served; a time (local at server or local at user device) the ad was rendered; a rendering attribute of the ad; a position of the ad within a Web page; a price that the advertiser will be charged for the impression; a price that the advertiser will be charged for a click; a price that the advertiser will be charged for a conversion; an identity of the server that chose the ad; search conditions that generated the page with which the ad was rendered; a next intermediate URL; a final destination URL; an identity of the ad creative; a topic of the content or other information about the content with which the ad was served (e.g., one or more concepts associated with the Web page on which the ad was displayed); an identity of the content with which the ad was served; information about other ads that were rendered along with the ad; a geolocation to which the ad was served; and user profile information of the user to which the ad was served, various other ad serving parameters, etc.

Ad properties may be encoded as a sequence of key-value pairs within the URL, where the key identifies the kind or type of property that the value is referring to.

Alternatively, properties can be encoded as a sequence of values within the URL, where an absolute or relative position of each of the values within the sequence is used to identify the kind of property that the value corresponds to. Both of these techniques can be combined. In one embodiment, such an intermediate URL with encoded ad properties complies with the syntax requirements specified in the paper, T. Berners-Lee, et al, "Uniform Resource Identifiers (URI): General Syntax," Network Working Group, Request for Comments: 2396 (August 1998) (incorporated herein by reference and referred to as "RFC 2396").

Exemplary methods for performing these operations are described in § 4.2.1 below. Exemplary apparatus are described in § 4.2.2 below. Finally, various refinements and alternatives are described in § 4.2.3 below.

§ 4.2.1 EXEMPLARY METHODS

Figure 4 is a flow diagram of an exemplary method 400 that may be used to encode ad property information in a manner consistent with the present invention. Property values (and property keys if they are included) may be represented as 5 characters, such as ASCII characters. These characters may have binary representations. These binary representations may be concatenated to define a sequence of bits. (Block 410) The sequence of bits may then be encoded into a sequence of K-nary characters. (Block 420) The generated sequence of K-nary characters may then be provided in a click URL (Block 430) before the exemplary 10 method 400 is left (Node 440).

Referring back to 420, K may represent a number of legal characters. For example, if legal characters for a click URL are limited to upper and lower alpha characters, digits, and the marks “-” “_” “.” “!” “~” “*” “” “(” and “)”, K may be 71 (=26+26+10+9). The encoding scheme may use a subset of the legal characters, but 15 this may cause the encoded information to be longer. Property keys and values might be in the form of numbers or sequences of ASCII characters. These numbers or sequences of ASCII characters can be encoded into a legal URL in such a way that they can be decoded.

Figure 5 is a flow diagram of an exemplary method 500 that may be used to 20 encode information, such as variable length information, (Recall, e.g., block 420 of Figure 4.) in a manner consistent with the present invention. The sequence of bits (Recall block 410 of Figure 4.) can be thought of as a large number, N. An enumerated list of characters that are legal within a URL is chosen for encoding and decoding. If there are K legal characters in the list, then N can be represented by $N \log K$ 25 characters. N is divided by K to obtain a result and a remainder. (Block 510) Since the remainder will be a number between 0 and K-1, it may be used to choose a legal character from the enumerated list of legal characters. (Block 520) The chosen legal character may then be used in the click URL. (Block 530) The whole number result of the division may then be repeatedly divided by K (Block 540), the sequence of 30 remainders is used to choose a sequence of characters from the K legal characters (Block 550), and the chosen characters are used in the click URL (Block 560). As

indicated by decision block 570 and block 580, this process is repeated if the whole number result is not less than K. If the whole number result is less than K, the exemplary method 500 is left (Node 590).

5 **§ 4.2.2 EXEMPLARY APPARATUS**

Figure 6 is high-level block diagram of a machine 600 that may be used to perform one or more of the operations discussed above. The machine 600 may include one or more processors 610, one or more input/output interface units 630, one or more storage devices 620, and one or more system buses and/or networks 640 for facilitating 10 the communication of information among the coupled elements. One or more input devices 632 and one or more output devices 634 may be coupled with the one or more input/output interfaces 630.

The one or more processors 610 may execute machine-executable instructions (e.g., C or C++ running on the Solaris operating system available from Sun 15 Microsystems Inc. of Palo Alto, California or the Linux operating system widely available from a number of vendors such as Red Hat, Inc. of Durham, North Carolina) to effect one or more aspects of the present invention. At least a portion of the machine executable instructions may be stored (temporarily or more permanently) on the one or more storage devices 620 and/or may be received from an external source via one or 20 more input interface units 630.

In one embodiment, the machine 600 may be one or more conventional personal computers. In this case, the processing units 610 may be one or more microprocessors. The bus 640 may include a system bus. The storage devices 620 may include system memory, such as read only memory (ROM) and/or random access 25 memory (RAM). The storage devices 620 may also include a hard disk drive for reading from and writing to a hard disk, a magnetic disk drive for reading from or writing to a (e.g., removable) magnetic disk, and an optical disk drive for reading from or writing to a removable (magneto-) optical disk such as a compact disk or other (magneto-) optical media.

30 A user may enter commands and information into the personal computer through input devices 632, such as a keyboard and pointing device (e.g., a mouse) for example.

Other input devices such as a microphone, a joystick, a game pad, a satellite dish, a scanner, or the like, may also (or alternatively) be included. These and other input devices are often connected to the processing unit(s) 610 through an appropriate interface 630 coupled to the system bus 640. The output devices 634 may include a

5 monitor or other type of display device, which may also be connected to the system bus 640 via an appropriate interface. In addition to (or instead of) the monitor, the personal computer may include other (peripheral) output devices (not shown), such as speakers and printers for example.

The ad server 210, user device (client) 250, search engine 220, content server
10 230, and/or e-mail server 240 may be implemented as one or more machines 600.

§ 4.2.3 REFINEMENTS AND ALTERNATIVES

To ensure that the properties encoded in the URL are genuine, a digital signature can be created from the property values, and/or from part of or the entire encoded URL.

15 If a private key is used to help generate the digital signature, then the author of the URL and properties can be verified. The digital signature may be generated using algorithms such as HMAC or other algorithms known in the art. The digital signature can be encoded within the URL just like any other property.

In a refined embodiment of the present invention, 64 enumerated characters may be provided (i.e., K=64) that are legal within a URL. This technique results in removing an integral number of bits from the binary representation of N and is therefore efficient to implement on modern microprocessors. Alternatively, K may be 2^x , and preferably less than 72. Thus, for example, K may be 2, 4, 8, 16, 32 and 64.

In some embodiments of the present invention, certain characters may be excluded from the set of K legal characters. For example, in one embodiment, one or more of the characters "control" "space" "<" ">" and "%" are excluded. In another embodiment, one or more of the characters "{" "}" "|" "\\" "\^" "[" "]" and "\`" are excluded.

In certain embodiments of the present invention, numbers may be represented by fewer bits using variable length encoding techniques. One technique is to represent a number as a sequence of shorter fixed bit-length numbers. For example, these shorter numbers may be one byte each. One bit within the shorter fixed-length number

may be used as a continuation bit. The other bits may be used to represent value bits for the original number. Value bits from the fixed bit-length numbers may be accumulated until a fixed bit-length number whose continuation bit is reset is encountered. The value bits may then be combined to form the number.

5

§ 4.3 CONCLUSIONS

As can be appreciated from the foregoing disclosure, the present invention can be used to help track ad selections and/or ad performance by encoding ad properties into a click URL.